



Argonne  
NATIONAL  
LABORATORY

*... for a brighter future*

**ALCF**

*Argonne Leadership  
Computing Facility*



U.S. Department  
of Energy

UChicago ►  
Argonne<sub>LLC</sub>



A U.S. Department of Energy laboratory  
managed by UChicago Argonne, LLC



# Argonne Leadership Computing Facility

## Overview

*Ray Bair, Director  
Argonne National Laboratory and  
University of Chicago*

*June 20, 2007*


# Over 20 years of Advanced Systems for DOE and Others

## ■ ACRF period [1983-1992]

- DOE's founding ACRF
- Explored many parallel architectures, developed programming models and tools, trained >1000 people

## ■ HPCRC period [1992-1999]

- Production-oriented parallel computing for Grand Challenges in addition to Computer Science.
- Fielded 1<sup>st</sup> IBM SP in DOE



## ■ TeraGrid [2001-present]

- Overall Project Lead
- Defining, deploying and operating the integrated national cyberinfrastructure for NSF
- 9 sites, 22 systems, 200TF

## ■ LCRC [2003-present]

- Lab-wide production supercomputer service
- All research divisions, 56 projects, 380 users

## ■ BlueGene Evaluation [2005-present]

- Founded BlueGene Consortium with IBM
  - 67 institutions, >260 members
  - Applications Workshop Series
  - Systems Software Collaborations

# ***DOE Leadership Computing Facility (LCF) Strategy***

- DOE SC selected the ORNL, ANL and PNNL teams (May 12, 2004) based on a competitive peer review
  - ORNL will deploy a series of systems based on Cray's XT3/4 architectures @ 250TF/s in FY07 and 1000TF/s in FY08/09
  - ANL will develop a series of systems based on IBM's BlueGene @ 100TF/s in FY07 and 250-500TF/s in FY08/FY09 with IBM's Next Generation Blue Gene
  - PNNL will contribute software technology
- DOE SC will make these systems available as capability platforms to the broad national community via competitive awards (e.g. INCITE Allocations)
  - Each facility will target ~20 large-scale production applications teams
  - Each facility will also support development users
- DOE's LCFs complement existing and planned production resources at NERSC
  - Capability runs will be migrated to the LCFs, improving NERSC throughput
  - NERSC will play an important role in training and new user identification

# ***Mission and Vision for the ALCF***

## **Our Mission**

Provide the computational science community with a world leading computing capability dedicated to breakthrough science and engineering.

## **Our Vision**

A world class center for computation driven scientific discovery that has:

- outstandingly talented people,
- the best collaborations with computer science and applied mathematics,
- the most capable and interesting computers and,
- a true spirit of adventure.

See <http://www.alcf.anl.gov/> for additional information

# *Decision to choose Blue Gene is Supported by*

- Blue Gene has been fielded within a factor of 3 of PF goal
  - *No other system is close to this scale (lower risk to scale to PF)*
- Applications community has reacted positively, though the set of codes is still limited, but is larger than expected, and some applications are doing extremely well
  - *For those applications that can make the transition, the BG platform provides outstanding scientific opportunity - many can, some can't*
- Blue Gene has been remarkably reliable at scale
  - *The overall reliability/TF appears to be at least an order of magnitude better than other platforms for which we have data*
- Power consumption is 2x better than other platforms
  - *Lower cost of ownership and window to the future of lower power*
- System Cost
  - *The cost of deploying a balanced system is lower than other platforms*

# Example Blue Gene Science Applications

- **Qbox** — FPMD solving Kohn-Sham equations, strong scaling on problem of 1000 molybdenum atoms with 12,000 electrons (86% parallel efficiency on 32K cpus @ SC05), achieved 207 TFs recently on BG/L
- **ddcMD** — many-body quantum interaction potentials (MGPT), 1/2 billion atom simulation, 128K cpus, achieved > 107 TFs on BG/L via fused dgemm and ddot
- **BlueMatter** — scalable biomolecular MD with Lennard-Jones 12-6, P3ME and Ewald, replica-exchange 256 replicas on 8K cpus, strong scaling to 8 atoms/node
- **GAMESS** — *ab initio* electronic structure code, wide range of methods, used for energetics, spectra, reaction paths and some dynamics, scales  $O(N^5-N^7)$  in number of electrons, uses DDI for communication and pseudo-shared memory, runs to 32,000 cpus
- **FLASH3** — produced largest weakly- compressible, homogeneous isotropic turbulence simulation to date on BG/L, excellent weak scaling, 72 million files 156 TB of data

QuickTime™ and a  
TIFF (LZW) decompressor  
are needed to see this picture.

QuickTime™ and a  
TIFF (LZW) decompressor  
are needed to see this picture.

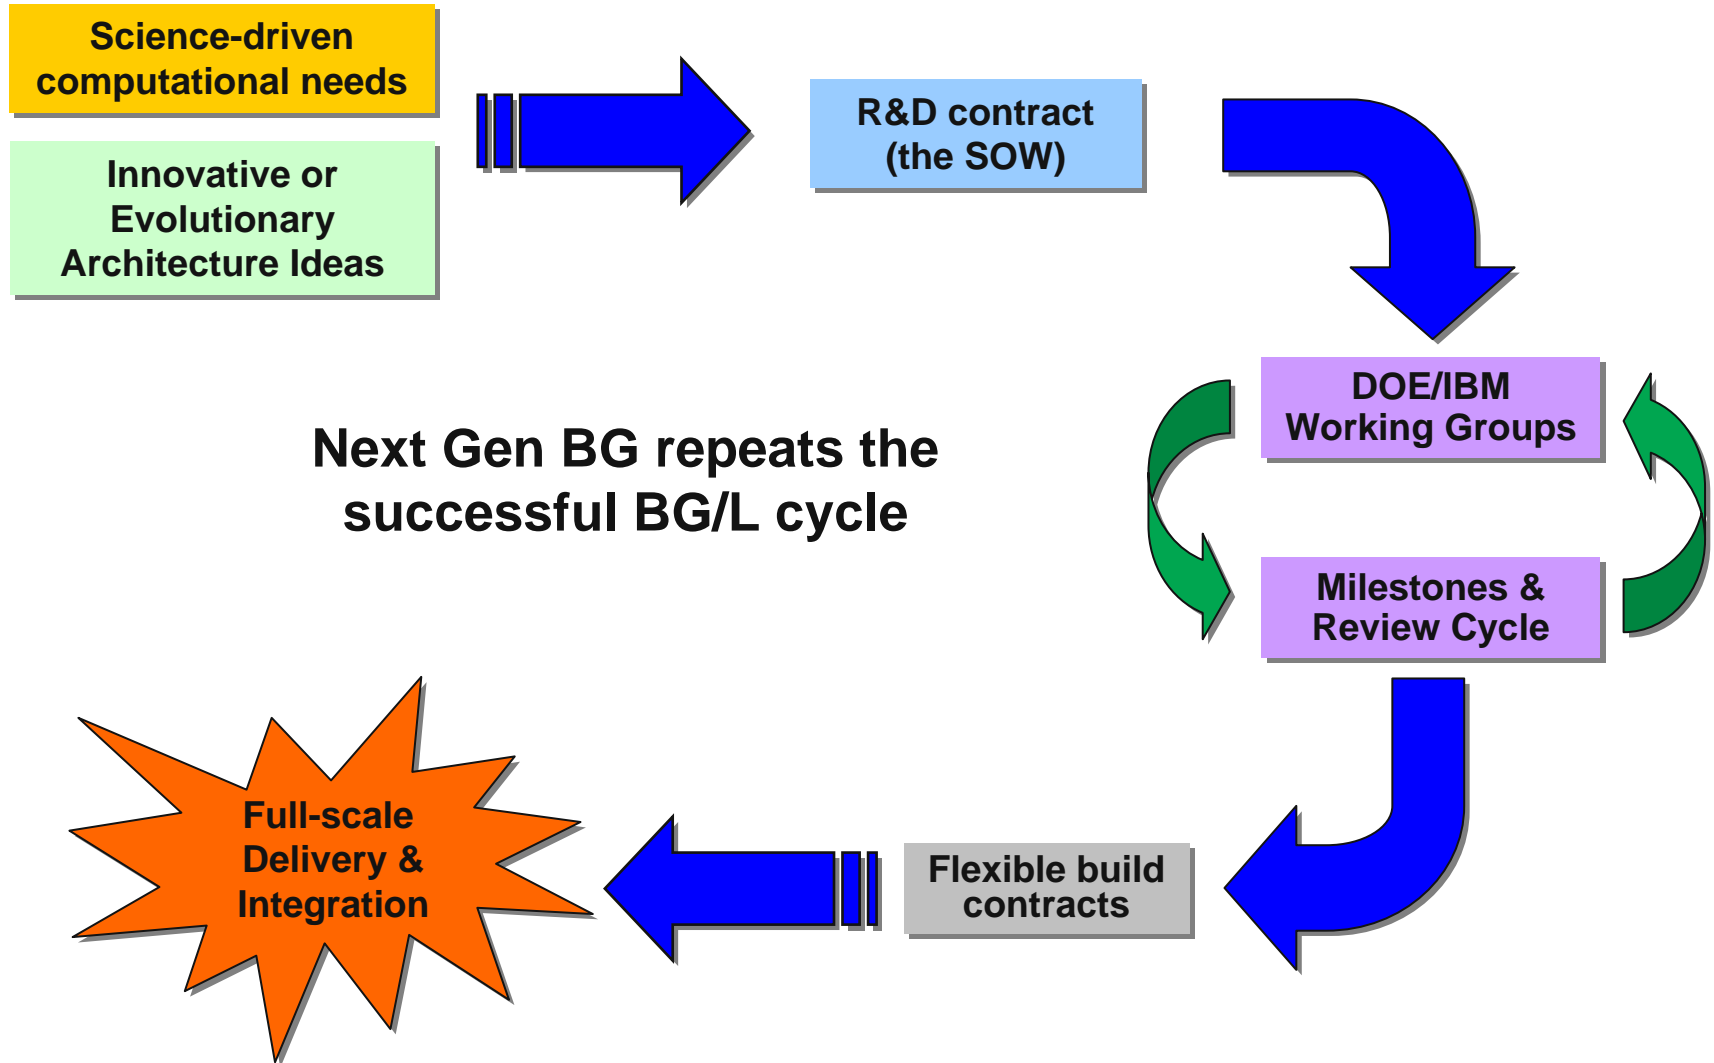
QuickTime™ and a  
TIFF (LZW) decompressor  
are needed to see this picture.

QuickTime™ and a  
TIFF (LZW) decompressor  
are needed to see this picture.



# DOE / IBM Partnership for Petascale Architectures

DOE NNSA/LLNL, DOE SC/ANL and IBM/Rochester/Watson



# Argonne Leadership Computing Facility

*Established 2006. Dedicated to breakthrough science and engineering.*

- **Current Blue Gene/L Capabilities**
  - BGL: 1024 nodes, 2048 cores, 5.7 TF, 512GB
  - Supports development + INCITE
- **Additional 2007 INCITE time at IBM T.J. Watson Research**
  - BGW: 20,480 nodes, 40,960 cores, 114 TF, 10 TB
- **Coming for early 2008 production**
  - 100+ TF next gen. system
  - Fast PB file system
  - Many PB tape archive
- **Then for early 2009 production**
  - Major Blue Gene upgrade
  - Next gen. file system



BGL



BGW

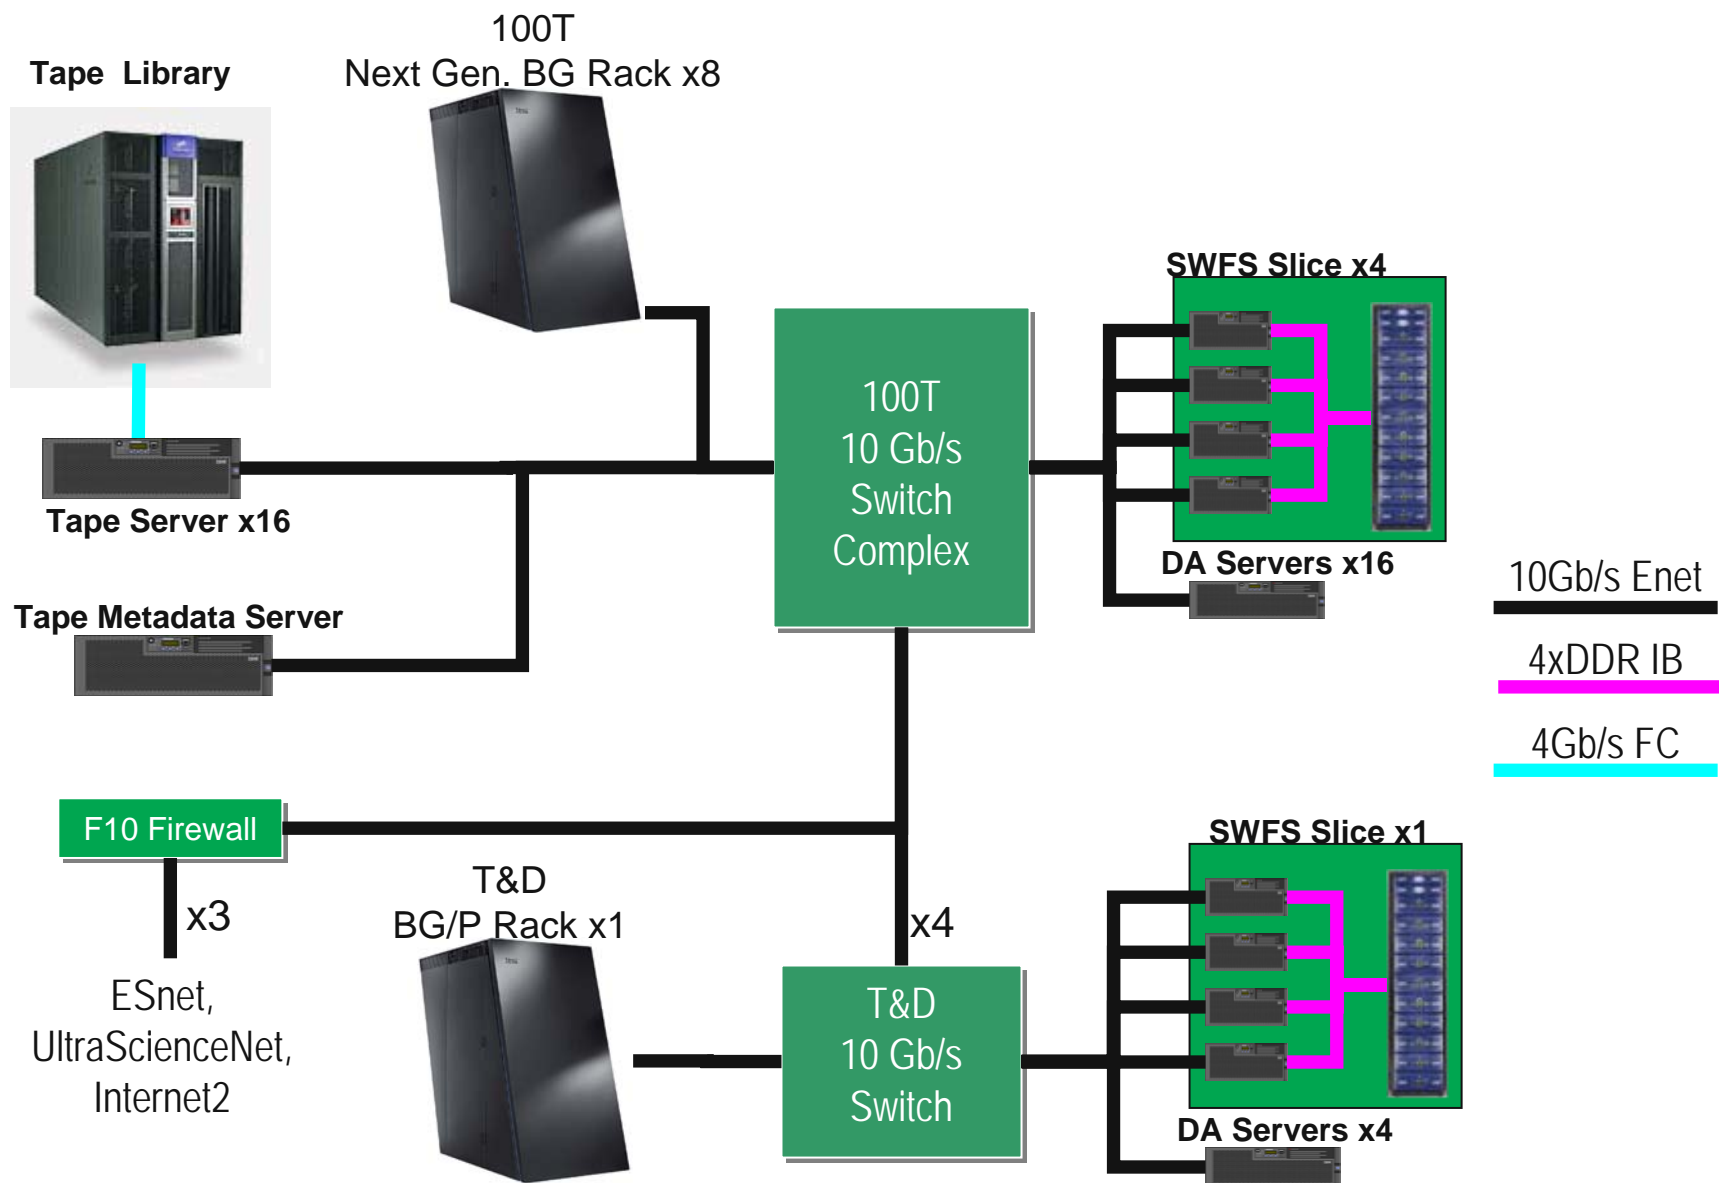


Next Generation Blue Gene



# ALCF Q3 FY2008

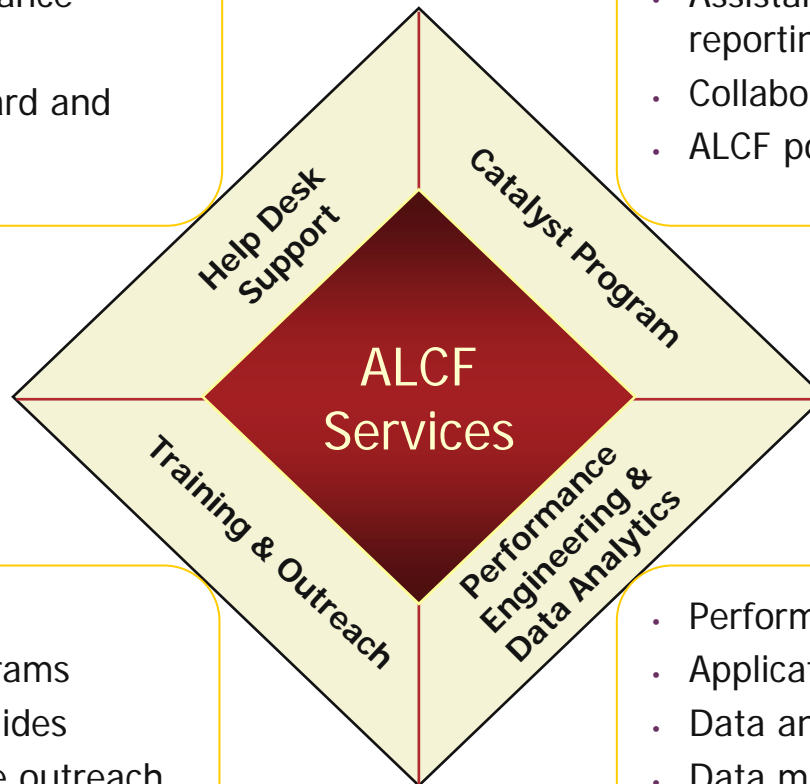
*At start of INCITE Production*



# ALCF Service Offerings

- Startup assistance
- User administration assistance
- Job management services
- Technical support (Standard and Emergency)

- ALCF project management
- Assistance with proposals, planning, reporting
- Collaboration within science domains
- ALCF point of coordination



- Workshops & seminars
- Customized training programs
- On-line content & user guides
- Educational and corporate outreach programs

- Performance engineering
- Application tuning
- Data analytics
- Data management services

# Catalyst Program

---

## ***What is it?***

ALCF Program to establish strategic collaborations with our leading project partners to maximize benefits from the use of various ALCF resources

---

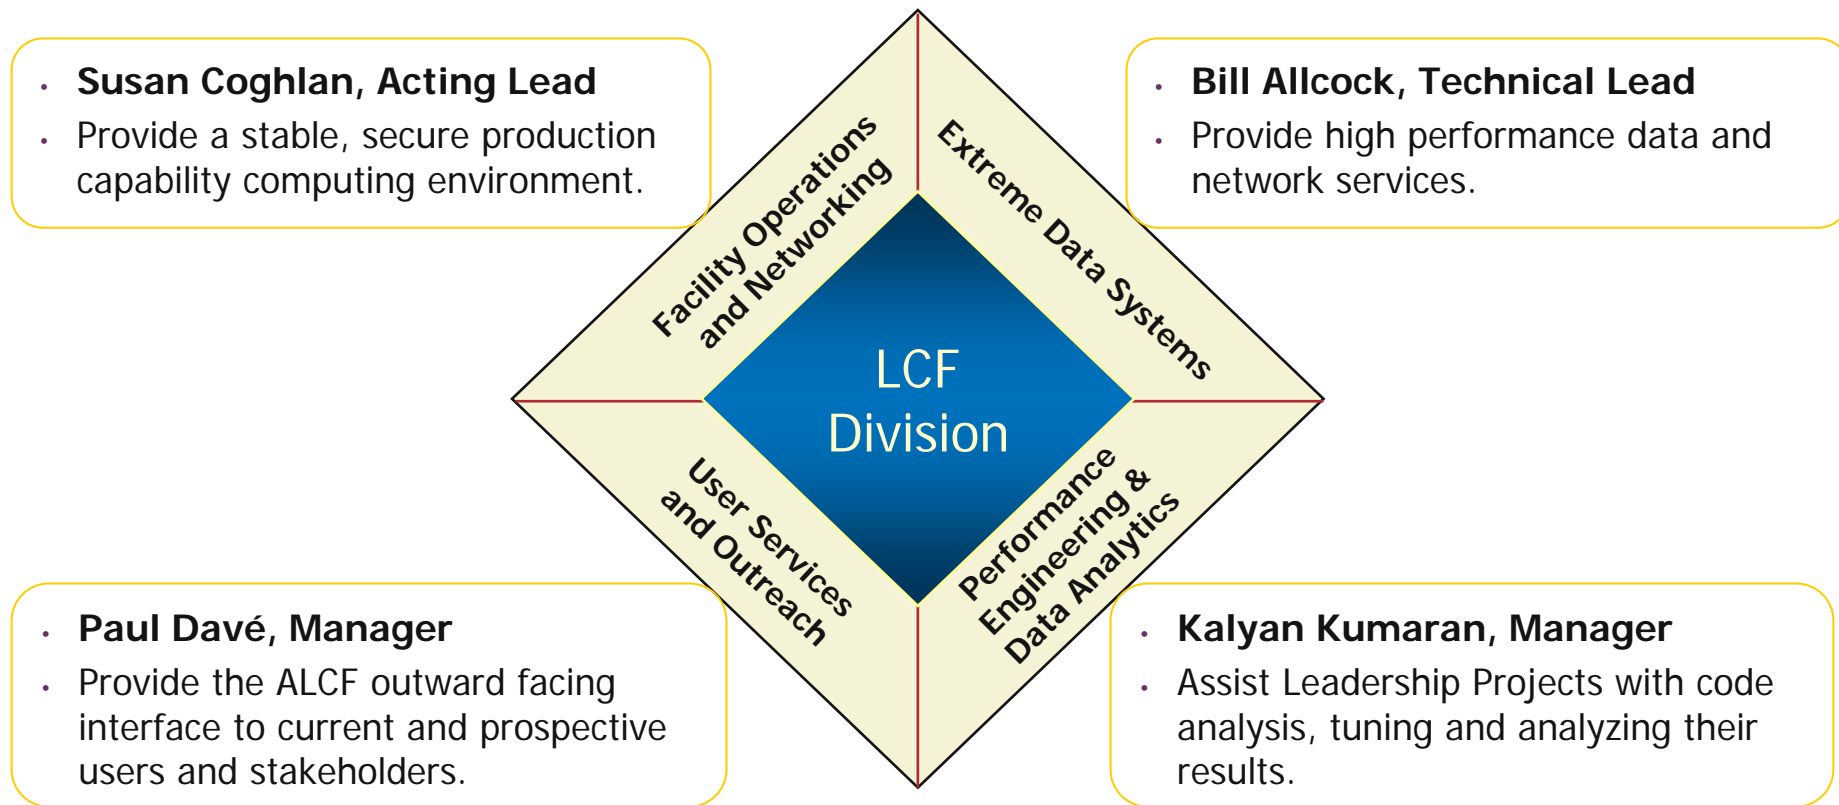
## ***What are the program features?***

- Full project lifecycle approach to collaboration with our project partners
  - Provide value-added services and support in conjunction with ALCF HW and SW resources
  - Tailor services-delivery program for the unique requirements of each research initiative
  - Maintain close contact with research teams through ongoing interactions with assigned ALCF Project Coordinator
- 

## ***What are the key objectives?***

- 'Jump-start' the use of ALCF resources for major ALCF projects
- Align availability of ALCF services and compute resources with the needs of researchers through joint project planning based on research goals and timing objectives
- Establish a spirit of collaboration to maximize the value that ALCF can bring to our project partners

# ALCF Organization



# ALCF Timeline

## 2004

- Formed of the Blue Gene Consortium with IBM
- DOE-SC selected the ORNL, ANL and PNNL teams for Leadership Computing Facility award

## 2005

- Installed 5 teraflops Blue Gene/L for evaluation

## 2006

- Began production support of 6 INCITE projects, with BGW
- Joined IBM and LLNL to design and develop next Blue Genes
- “Lehman” Peer Review of ALCF campaign plans

## 2007

- Increased to 9 INCITE projects; continued development projects
- Install 100 teraflops next gen Blue Gene system (late 2007)

## 2008

- Begin support of INCITE projects on next gen Blue Gene
- Add 250-500T teraflops Blue Gene system